arXiv:1605.09772v1 [cs.SY] 31 May 2016

Technical Report: Directed Controller Synthesis of Discrete Event Systems

Daniel Ciolek[1], Victor Braberman[1,3], Nicolás D'Ippolito[1,2,3] and Sebastián Uchitel[1,2,3]

***Abstract*— This paper presents a Directed Controller Synthesis (DCS) technique for discrete event systems. The DCS method explores the solution space for reactive controllers guided by a domain-independent heuristic. The heuristic is derived from an efficient abstraction of the environment based on the componentized way in which complex environments are described. Then by building the composition of the components on-the-fly DCS obtains a solution by exploring a reduced portion of the state space. This work focuses on untimed discrete event systems with safety and co-safety (i.e. reachability) goals. An evaluation for the technique is presented comparing it to other well-known approaches to controller synthesis (based on symbolic representation and compositional analyses).**

I. Introduction

Discrete Event Systems (DES) are dynamic systems that react to the occurrence of diverse discrete events. DES arise in many domains including robotics, logistics, manufacturing, and communication networks. These applications require control and coordination to ensure that the system goals are achieved. Numerous techniques for automatically verifying the correctness of controllers have been devised in the area of model checking. A different yet related approach, Controller Synthesis, pursues the automatic construction of controllers satisfying a formal specification.

The field of synthesis of controllers for DES was introduced by Ramadge and Wonham [1] for controlling systems within a given set of constraints. In this setting, the environment (also called plant) and goals are specified using a formal language, and a procedure generates a correct by construction controller (or supervisor). The environment is usually modeled with state machines whose event sets are partitioned into controllable and uncontrollable actions. A controller must achieve its goals by dynamically disabling some of the controllable actions.

An alternative formulation of the controller synthesis problem considers the setting of ω-regular languages [2], instead of an automata-based DES. In this setting two-player games with Boolean objectives (where a win of one player coincides with a loss by the other player) have been suggested to reason about the interaction between the program and the environment [3], [4].

We address control problems for behavior models expressed as deterministic Labeled Transition Systems (LTS) and parallel composition ($\|$) defined broadly as a synchronous product. This setting allows to ease the modeling of complex environments by describing its behavior compositionally. Such descriptions are compact and, hence, obtaining the complete behavior of the environment via composition produces an exponential explosion. For this reason we want to avoid the computation of the composition ahead of time, building it on-the-fly instead (hopefully obtaining a solution by exploring a reduced portion of the state space).

Hereinafter we present the Directed Controller Synthesis (DCS) of DES for safety and co-safety (i.e. reachability) goals. Informally, we want the controller to *always* avoid safety violations while *eventually* ensuring co-safety objectives. More formally, we make use of the Linear Temporal Logic (LTL) [5] operators. Thus, given the model of the environment components $E_0, \ldots, E_n$, the safety goal G_S and the co-safety goal G_C, we look for a component M such that when composed with the environment it satisfies the goals, namely $E_0 \| \ldots \| E_n \| M \models \Box\neg G_S \wedge \Diamond G_C$.

DCS explores the solution space for reactive controllers on-the-fly guided by a domain-independent heuristic. The main contribution of this paper is the heuristic derived from an abstraction of the environment that exploits the componentized way in which complex environments are described. We highlight that componentization not only simplifies modeling, but also allows for preprocessing procedures that can improve the applicability of controller synthesis techniques.

The paper is organized as follows. Section II comments on previous and related work. Section III provides background. Section IV comments on details to take into account during the on-the-fly exploration of the state space. Section V presents the LTS abstraction and a heuristic derived from it. Section VI reports on the results obtained with our implementation. Finally, section VII discuses some conclusions and avenues for future work.

II. Related Work

Traditional controller synthesis techniques for safety objectives look for maximal (i.e. least restrictive) controllers [6], since safety properties can many times be trivially satisfied by restricting the system to an unproductive zone. Maximality comes at a cost in complexity given that it requires to explore the complete state space.

Monolithic approaches to controller synthesis work with the complete state space but, since it is exponential with respect to the size of the components, explicit representations are impractical. The Model Based Planner (MBP) presented in [7] uses Binary Decision Diagrams (BDD) [8] to keep a symbolic representation of the state space, in an attempt to cope with the state explosion problem. MBP is based on

1. Departamento de Computación, Universidad de Buenos Aires, Argentina
2. Department of Computing, Imperial College, London, UK
3. CONICET

This work was submitted to the 55^{th} IEEE Conference on Decision and Control and was partially supported by grants ERC PBM-FIMBSE, ANPCYT PICT 2011-1774, ANPCYT PICT 2012-0724, ANPCYT PICT 2013-2341, UBACYT 20020130100384BA, UBACYT 20020130300036BA, CONICET PIP 112 201301 00688 CO, PIDDEF 04/ESP/15/BAA.

the NuSMV model checker and hence, its input is a model described in the SMV language, implementing different variations of goals (including liveness with CTL).

Contrarily to monolithic approaches, compositional approaches to controller synthesis (as the one implemented by Supremica [9]) build a controller for each component, such that when composed achieve the system requirements. Although this allows to greatly reduce the impact of the state explosion problem it also restricts its application to simple goals, such as safety, since it is not possible to guarantee composability for richer goals.

Controller synthesis is also related to the area of Automated Planning, a branch of artificial intelligence. However, the work in planning has been oriented mainly towards non-reactive environments, which are insufficient in the setting in which controller synthesis is applied. MBP provides a translation from a planning input to an SMV script and shows that a planning problem can be solved as a control problem by taking advantage of the advances on model checking [10].

In planning the state space is usually represented explicitly, but the construction is performed on-the-fly guided by heuristics [11]. Informed search procedures – like A* [12] – are used to perform a goal-directed exploration, generally obtaining a solution by inspecting a reduced portion of the state space. Many variations of the planning problem have been attacked this way and some – especially those that contemplate nondeterminism such as FOND planning [13] – are closely related to controller synthesis.

Inspired by planning, informed search procedures have been introduced in model checking to accelerate the search for an error [14], [15]. Although with this approach a trace to an error can be found faster, verifying the correctness of a model still requires to explore the complete state space. Despite the positive results of these techniques, their application to controller synthesis has been scarcely explored.

In [16] an on-the-fly synthesis method is presented for discrete and dense-time system, using depth-first search *without a heuristic guide*. For linear hybrid systems, an on-the-fly informed search algorithm is presented in [17], which excels at design space exploration by quickly assessing whether a problem has no solution, *rather than guiding the exploration towards a goal*. Finally the Controller Synthesis Module of the CIRCA architecture [18], which works on a timed automaton model, uses heuristics with a *limited lookahead* and a verification procedure to detect and prune bad choices. In all these approaches the on-the-fly search procedure is similar to ours, but the setting differs since we consider as input the specification of the environment components, and hence *the abstractions and heuristics diverge*.

Herein we present an extension to our previous work on the Modal Transition System Analyser (MTSA) [19] tool that now includes the DCS algorithm. As mentioned before, in the present work we focus on safety and co-safety (i.e. reachability) goals, the former being optional while the latter mandatory since the heuristic needs a desirable end to properly direct the search.

III. Background

Definition 1 (Labeled Transition Systems)**.** A *Labeled Transition System* (LTS) is a tuple $E = (S, A, \rightarrow, s_0)$, where S is a finite set of states, A is its alphabet, $\rightarrow \subseteq (S \times A \times S)$ is a transition relation, and $s_0 \in S$ is the initial state.

Notation 1. Let E and M be LTSs such that:

- $E = (S_E, A_E, \rightarrow_E, s_0)$, with s and s' states of S_E
- $M = (T_M, A_M, \rightarrow_M, t_0)$, with t and t' states of T_M

Notation 2 (Step)**.** We denote $(s, l, s') \in \rightarrow$ by $s \xrightarrow{l} s'$.

Definition 2 (Parallel Composition)**.** The *Parallel Composition* ($\|$) is a symmetric operator such that it yields an LTS $E\|M = (S_E \times T_M, A_E \cup A_M, \rightarrow_{E\|M}, \langle s_0, t_0 \rangle)$, where $\rightarrow_{E\|M}$ is a relation that satisfies the following rules:

$$\frac{s \xrightarrow{l}_E s'}{\langle s,t\rangle \xrightarrow{l}_{E\|M} \langle s',t\rangle}\ l \in (A_E \setminus A_M) \qquad \frac{t \xrightarrow{l}_M t'}{\langle s,t\rangle \xrightarrow{l}_{E\|M} \langle s,t'\rangle}\ l \in (A_M \setminus A_E)$$

$$\frac{s \xrightarrow{l}_E s', t \xrightarrow{l}_M t'}{\langle s,t\rangle \xrightarrow{l}_{E\|M} \langle s',t'\rangle}\ l \in (A_E \cap A_M)$$

Example 1. *In Fig.1 we show an example of the application of the parallel composition operator ($\|$). Note that synchronizing actions produce an update in the states of both LTSs, while nonshared actions produce an update in only one LTS at a time. Observe also that shared actions that are not available from both LTSs at a given point cannot be executed (e.g. d available from t_0 but not from s_0 cannot be executed from the initial state $\langle s_0, t_0 \rangle$).*

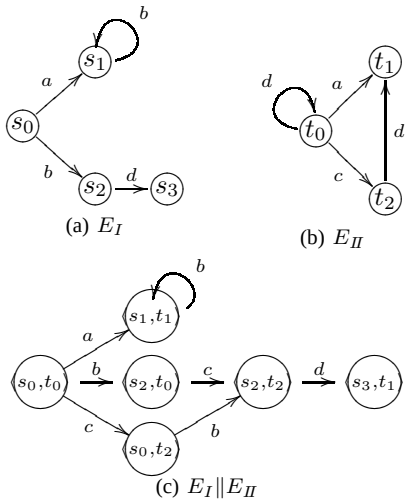

Fig. 1. Parallel Composition Example

Given a partition of the alphabet in controllable and uncontrollable actions, a controller M must achieve its goal G by disabling some of the controllable actions in E, while only monitoring uncontrollable actions. We say that a controller M satisfies a goal G in an environment E if and only if every trace accepted by $E\|M$ satisfies G.

Definition 3 (Trace)**.** A trace of an LTS E is a sequence of labels $\pi{=}l_0, l_1 \ldots$ of A_E, for which there exists a sequence of states $s_0, s_1, \ldots$ of S_E such that $\forall i{\geq}0 \;.\; s_i \xrightarrow{l_i}_E s_{i+1}$.

Definition 4 (Safety)**.** We say that a trace π satisfies a *safety* goal G_S given as a set of labels, if and only if a label of G_S never occurs in π.

Definition 5 (Co-Safety)**.** We say that a trace π satisfies a *co-safety* goal G_C given as a set of labels, if and only if there is at least one occurrence of a label of G_C in π.

Example 2. *In Fig.2 we show two controllers for $E_I \| E_{II}$ that are guaranteed to reach $\{d\}$ assuming that only a, b and c are controllable. The controller in Fig.2a simply disables a from the initial state to avoid entering a state where d is no longer reachable. Note that this controller does not force the order of actions b and c, that is, it imposes the least restrictive constraints. On the other hand, the controller in Fig.2b forces the order between actions b and c, hence, despite being correct it is not a maximal controller.*

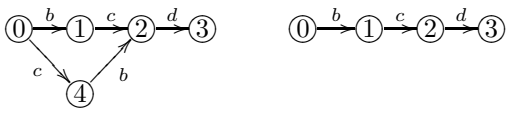

(a) Maximal (b) Nonmaximal

Fig. 2. Controllers for $E_I \| E_{II}$ reaching $\{d\}$

Notation 3. States of the composition are ordered pairs, however we may need to refer to the *step* relation for any ordering. For this purpose we may use the following notation:

$$\{s,t\} \xrightarrow{l}_{E\|M} \{s',t'\} \Leftrightarrow \langle s,t\rangle \xrightarrow{l}_{E\|M} \langle s',t'\rangle \vee \langle s,t\rangle \xrightarrow{l}_{E\|M} \langle t',s'\rangle \vee \langle t,s\rangle \xrightarrow{l}_{E\|M} \langle s',t'\rangle \vee \langle t,s\rangle \xrightarrow{l}_{E\|M} \langle t',s'\rangle$$

When using the set notation we relax the convention that s and s' belong to E while t and t' belong to M since the order of the states is made irrelevant.

Notation 4. For brevity we denote that a set of states $s_0, s_1, \ldots, s_n$ belongs to a set q, that is to say that $\forall i \;.\; 0 \leq i \leq n \Rightarrow s_i \in q$, as $[s_1, s_2, \ldots, s_n]_q$.

IV. On-the-fly Exploration

In this section we discuss the main algorithm that constitutes DCS. Given a heuristic function that estimates the distance from a state to a goal, DCS looks for a controller by computing the parallel composition on-the-fly guided by the heuristic. The procedure is a modification of Best First Search (a classical informed search procedure) adapted to account for uncontrollable actions. We keep a priority queue of *open* states, ordered by their estimated distance to a goal, initialized only with the initial state. At each iteration the algorithm takes the most promising state from the queue and expands a child state following its best unexplored action. The expanded state is evaluated using the heuristic, and it is then inserted in the *open* queue (the value of the state is set as the estimate of its best unexplored action).

For controllable states we need only one successful successor, while for uncontrollable states all their successors have to be able to reach a goal. Thus we allow a controllable state to remain in the *open* queue as long as it has unexplored actions to broaden the search (i.e. competing with its descendants). Mixed states, that is states with both controllable and uncontrollable actions, are treated as uncontrollable states. This is because mixed states represent a race condition between the controller and the environment, which can always be won by the environment. Thereby the validity of the controller cannot depend on the result of race conditions. This treatment of controllable and uncontrollable nodes is similar to AND/OR search procedures [20].

When a state is recognized either as a co-safety goal or a safety error, it is marked and this information is propagated back to its ancestors until an interrupting state is reached and reopened. The propagation of an error is interrupted when a controllable ancestor with unexplored actions is reached. Whereas, the propagation of a goal is interrupted when an uncontrollable ancestor with unexplored actions is reached. The process is repeated until the initial state is marked as a goal or an error. In the former the states reachable from the initial state form a controller, while in the latter there is no controller.

Observe that while there is an obvious advantage in exploring the most promising controllable action first, there is no such advantage in exploring the best uncontrollable action since all such actions must lead to a goal. Thus we could avoid the computation of the heuristic for uncontrollable states. Instead we use the reverse ranking and expand the least desirable uncontrollable action first. We do this in an attempt to find an error as fast as possible and consequently being able to close the state avoiding futile exploration.

In the worst case, the heuristic misguides the search and the algorithm explores the complete state space. With the additional cost of computing the heuristic for each state, the complexity of the algorithm is worse than its monolithic counterpart. That is, DCS could take an exponential amount of time. In spite of that, if the heuristic accurately guides the exploration with a small computational overhead, great savings could be obtained.

V. Abstraction and Heuristic

In this section we present the abstraction we use and the heuristic it induces. The main goal of the heuristic is to provide an estimate of the distance from a state to a goal without computing the parallel composition. However, in order to provide informative estimates the effects of synchronization need to be taken into account to some point. For this reason we build an abstraction of the environment that, instead of considering the cross product of states, it works on sets of states. More precisely, if the states contained in the set can synchronize (with the standard parallel composition), then the synchronizing action must be available from the set.

At first glance considering the power set of states may seem detrimental, but we also apply a rule of monotonic growth. That is, the reachable sets of states are restricted to

sets that contain all the traversed states from the initial state. Thus the abstraction size is polynomial with respect to the number of states in the components. We then use the length of the paths that reach a co-safety goal in the abstraction as an estimate of the real distance to the goal.

A. LTS Abstraction and Composition

Intuitively, we build an abstraction that, once it reaches a point where a state is covered, it never drops a state. The abstraction behaves as if, after a transition, not only a new state is covered but also the source state is not left. This progressively distances the abstraction from the real environment (as it collects states), yet it keeps information about the causal relationship between actions.

We begin by defining the *LTS abstraction* for a single LTS, that yields an LTS whose reachable region forms *a sequence of sets of states*.

Definition 6 (LTS Abstraction)**.** The abstraction of an LTS E is an LTS $\tilde{E} = (2^{S_E}, 2^{A_E}, \rightsquigarrow_E, \{s_0\})$, where $\rightsquigarrow_E$ is the relation that satisfies:

$$q \overset{A}{\rightsquigarrow_E} q' \Leftrightarrow \big(\forall s, s', a \,.\, [s]_q \wedge s \xrightarrow{a}_E s' \Rightarrow [s, s']_{q'} \wedge a \in A\big) \wedge \big(\forall s' \,.\, [s']_{q'} \Rightarrow [s']_q \vee \exists s, a \,.\, [s]_q \wedge a \in A \wedge s \xrightarrow{a}_E s'\big)$$

Observe that in the abstraction not only the states are sets of the original states, but also the labels are sets of the original labels. Furthermore, there is only one set of actions available from each set of states, that is, the maximal set of enabled actions.

We now define the *abstracting composition*, an operation for composing these abstracted LTSs. Informally, this operation works by applying a relaxed synchronization rule that follows that of ($\|$), but considering sets of states.

Definition 7 (Abstracting Composition)**.** The *Abstracting Composition* ($\S\!\!\S$) is a symmetric operator in the spirit of ($\|$) that yields an abstract LTS $\tilde{E} \S\!\!\S \tilde{M} = (S_{\tilde{E}} \cup S_{\tilde{M}}, A_{\tilde{E}} \cup A_{\tilde{M}}, \rightsquigarrow_{\tilde{E}\S\!\!\S\tilde{M}}, \tilde{m}_0 \cup \tilde{t}_0\})$ and $\rightsquigarrow_{\tilde{E}\S\!\!\S\tilde{M}}$ is the relation that satisfies:

$$q \overset{A}{\rightsquigarrow_{\tilde{E}\S\!\!\S\tilde{M}}} q' \Leftrightarrow \big(\forall s,s',t,t',a \,.\, [s,t]_q \wedge \langle s,t\rangle \xrightarrow{a}_{E\|M} \langle s',t'\rangle \Rightarrow [s',s,t',t]_{q'} \wedge a \in A\big) \wedge \big(\forall s' \,.\, [s']_{q'} \Rightarrow [s']_q \vee \exists s,t,t',a \,.\, [s,t]_q \wedge a \in A \wedge \{s,t\} \xrightarrow{a}_{E\|M} \{s',t'\}\big)$$

Note that we do not need to compute $E\|M$ since we only check locally that states s and t synchronize following the rules in Def.2.

Property 1. Given LTSs E and M the number of sets of states reachable in $\tilde{E}\S\!\!\S\tilde{M}$ is, in worst case, $|S_E| + |S_M|$.

Proof. The proof is immediate given that once every state in $S_M \cup S_E$ is included in a set of states q, no other set of states can be reached. Thus, in the worst case each step contributes only one fresh state, leading to $|S_E|+|S_M|$ steps before reaching the last possible fresh state. □

Example 3. *In Fig.3 we show examples for the LTS abstraction (Def. 6) and the abstracting composition (Def. 7) of LTSs E_I and E_{II}. In Fig.3a the abstraction of E_I is depicted. The first step of $\tilde{E}_I$ is constituted by the actions a and b (the actions available from s_0), reaching a set of states with s_1 (the target of a) and s_2 (the target of b), plus states already reached in previous steps (i.e. s_0). The second step only adds the action d available from s_2 reaching s_3, once all the reachable states are included in the set a self loop captures potential infinite behaviours.*

In Fig.3b a similar construction is shown for the abstracting composition between $\tilde{E}_I$ and $\tilde{E}_{II}$. Note that despite the fact that d is available from t_0 it is not available from $\{s_0, t_0\}$ since it is a synchronizing action not enabled at s_0.

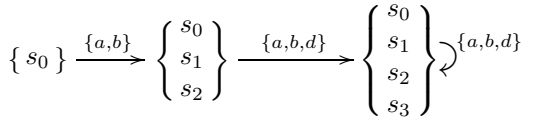

(a) $\tilde{E}_I$

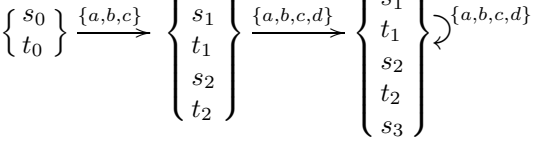

(b) $\tilde{E}_I \S\!\!\S \tilde{E}_{II}$

Fig. 3. LTS abstraction and composition examples

Despite looking very dissimilar there is a strong relation between a model and its abstraction. In particular if π is a trace of $E\|M$, π is a path contained in $\tilde{E}\S\!\!\S\tilde{M}$. We say that a sequence of actions $\pi{=}l_0, l_1 \ldots$ is a *contained path* in $\tilde{E}\S\!\!\S\tilde{M}$ if and only if there is a trace $\Pi{=}L_0, L_1 \ldots$ of $\tilde{E}\S\!\!\S\tilde{M}$ such that $\forall i \,.\, l_i \in L_i$. Moreover, we may abuse notation and denote that for a given label l_i exists a set L_i such that $l_i \in L_i \wedge q \overset{L_i}{\rightsquigarrow}_{\tilde{E}\S\!\!\S\tilde{M}} q'$ using directly l_i as the label of the transition as follows: $q \overset{l_i}{\rightsquigarrow}_{\tilde{E}\S\!\!\S\tilde{M}} q'$.

Note that by Def. 7 every trace of $E\|M$ is a path contained in $\tilde{E}\S\!\!\S\tilde{M}$. However, not every contained path can be mapped to a trace in $E\|M$. Therefore after computing the abstraction some paths will relate to traces while others will not, yet a priori we have no way to distinguish between these cases. We can easily filter some paths that are obviously uninteresting, such as the paths that contain noncomponent-steps or non-synchronizing-steps (e.g. $s_0 \xrightarrow{c}_{\tilde{E}_I \S\!\!\S \tilde{E}_{II}} t_2$). For this reason we define *abstracting path* to more precisely capture the notion of path that will help us estimate the distance to a goal.

Definition 8 (Abstracting Path)**.** We say that a sequence of actions $l_0, l_1, \ldots$ is an *abstracting path* of $\tilde{E}\S\!\!\S\tilde{M}$ if and only if there is a sequence of states $s_0, s_1, \ldots$ of $S_E \cup S_M$ and a sequence of sets of states $q_0, q_1, \ldots$ of $S_{\tilde{E}\S\!\!\S\tilde{M}}$ such that:

$$\forall i \,.\, ([s_i]_{q_i} \wedge q_i \overset{l_i}{\rightsquigarrow}_{\tilde{E}\S\!\!\S\tilde{M}} q_{i+1} \wedge \exists t,t' \,.\, \{s_i, t\} \xrightarrow{l_i}_{E\|M} \{s_{i+1}, t'\}) \vee (s_i{=}s_{i+1} \wedge l_i{=}\tau)$$

This definition allows not only for valid intra-LTS steps (e.g. $s_0 \xrightarrow{a}_{\tilde{E}_I \S\!\!\S \tilde{E}_{II}} s_1$), but it also allows for inter-LTS steps as long as they are potential synchronization steps (e.g. $s_0 \xrightarrow{a}_{\tilde{E}_I \S\!\!\S \tilde{E}_{II}} t_1$). This is important since an isolated component may not be able to reach a goal, but it may be necessary for an intermediate synchronization step. Observe that we add a

special action τ to "delay" a step from a state $[s_i]_{q_i}$ to actions available from $[s_i]_{q_{i+1}}$. This allows to consider paths that skip necessary synchronization steps, while still considering them in the estimation. This definition filters the paths that are not solely constituted by these intra-LTS and inter-LTS steps. However, it remains that not every abstracting path is a trace, yet we use the length of these paths as an estimate of the expected distance to a co-safety goal.

Property 2. Given LTSs E and M every trace π of $E\|M$ is an *abstracting path* of $\tilde{E}\S\tilde{M}$.

Proof. The proof is straightforward, since given that $\pi = l_0, l_1, \ldots$ is a trace of $E\|M$ there exists a sequence of states:

$$\langle s_0, t_0\rangle \xrightarrow{l_0}_{E\|M} \langle s_1, t_1\rangle \xrightarrow{l_1}_{E\|M} \ldots$$

Then by Def. 7 there exists a sequence of sets of states $q_0, q_1, \ldots$ such that $\forall i \,.\, [s_i, t_i]_{q_i}$ and

$$q_0 \overset{l_0}{\rightsquigarrow}_{\tilde{E}\S\tilde{M}} q_1 \overset{l_1}{\rightsquigarrow}_{\tilde{E}\S\tilde{M}} \ldots$$

Therefore by Def. 8 π is an *abstracting path* of $\tilde{E}\S\tilde{M}$. □

We compute the estimate simply by taking the length of an *abstracting path* that reaches a co-safety goal. Paths that reach a safety violation are considered to have an infinite (∞) distance to the goal. Interestingly, this can never overestimate the distance to the goal, but it can underestimate it. Heuristics that do not overestimate the distance to the goal are called admissible and have been studied in the literature since they enjoy useful properties (e.g. early error detection).

Lemma 1 (Admissibility)**.** The length of an *abstracting path* that reaches a co-safety goal is an admissible heuristic, that is, it does not overestimate the distance to the goal.

Proof. Note that in the worst case a path π reaching a goal is a trace of the environment, thus its length estimates exactly the distance to the goal. In any other case π is not a valid sequence of actions of the environment, that is to say is an artifact of the weakened composition rules of the abstraction. Hence π may skip some necessary intermediate steps required to actually reach the goal in the environment. Therefore, the length of π underestimates the distance to the goal. □

B. Heuristic Computation

In this section we show how we build the abstraction and compute the heuristic estimates for the actions available from a given state. For computing the heuristic we need all the *abstracting paths* reaching a goal or an error from an initial state. For this we build a graph whose paths are all such *abstracting paths*, which is naturally induced from Def. 8. However, instead of explicitly representing the special action τ, we represent the "distance" between states by a positive integer that we call *generation*.

Definition 9 (Generation)**.** Given an abstracted LTS $\tilde{E}$ we define the *generation* of a state s of $S_{\tilde{E}}$ as the index g of the *sequence of sets of states* of $\tilde{E}$ in which the state s appears for the first time, denoted s^g. More formally,

$$s^g \Leftrightarrow g = \max_i \{\, i \mid \forall\, 0 \le j < i, \exists l_j \,.\, q_j \overset{l_j}{\rightsquigarrow}_{\tilde{E}} q_{j+1} \wedge s \notin q_j\}$$

The *generation* is useful since given a transition from a state s_i into s_j with generations n and m respectively, we can deduce the "distance" between them by subtracting $|m-n|$.

Definition 10 (Abstracting Path Graph)**.** Given LTSs E and M, the *Abstracting Path Graph* of $\tilde{E}\S\tilde{M}$ is a graph $\mathcal{G}(\tilde{E}\S\tilde{M}) = (\mathcal{V}, \mathcal{E}, gen)$ where:

- $\mathcal{V} = S_E \cup S_M$, is a set of vertices formed by states
- $\mathcal{E} = \{(s_i, l, s_j) \mid \exists q, q', t, t' \,.\, [s_i, t]_q \wedge q \overset{l}{\rightsquigarrow}_{\tilde{E}\S\tilde{M}} q' \wedge \{s_i, t\} \xrightarrow{l}_{E\|M} \{s_j, t'\}\}$, is a set of labeled edges
- $gen : V \to \mathbb{N}$, is a function mapping states to their respective *generations*

Informally, two vertices of the graph are connected by an edge if and only if there is a step, connecting the corresponding states, that is part of an *abstracting path*.

Example 4. *In Fig.4 we show the abstracting path graph of $\tilde{E}_I\S\tilde{E}_{II}$. Each edge of the graph represents a possible step in the abstraction that is a part of an abstracting path. The first column of the graph contains the states at generation 0 (s_0 and t_0), the second column contains the fresh states at generation 1 (s_1, s_2, t_1 and t_2), while the third column contains the fresh states at generation 3 (s_3). An edge that skips one or more columns (or generations) represents a fragment of an abstracting path with τ-steps. For example, the edge labeled d connecting t_0 with s_3 represents a τ-step followed by action d. This is because d is not available from the initial state and is delayed (i.e. the path has a length of 2).*

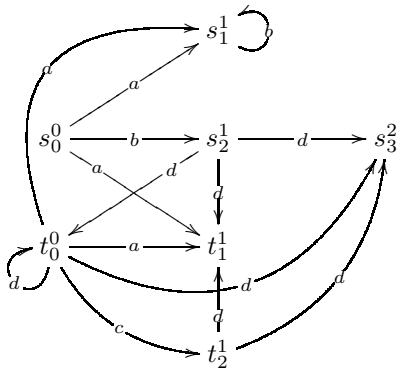

Fig. 4. Abstracting Path Graph $\mathcal{G}(\tilde{E}_I\S\tilde{E}_{II})$

We can use this graph to build a ranking of the available actions from the initial state of $E_I\|E_{II}$, namely $\langle s_0, t_0\rangle$, by measuring the distance to the co-safety goal $\{d\}$. The available actions are a, b and c (as mentioned before d is disabled from the initial state). The shortest paths to d starting with b and c have a length of 2, while there are no paths reaching d starting with a. Hence, we estimate an infinite distance to the goal for every path starting with a. Ranking the actions by these estimates directs us to explore the actions b and c. In spite of that, we are underestimating the distance to the goal since our estimation assumes that immediately after either b or c we would be able to execute d. This is an artifact of the weak composition rules since both b and c need to be executed (in any order) before d. Still, these estimates effectively guide the exploration towards the goal.

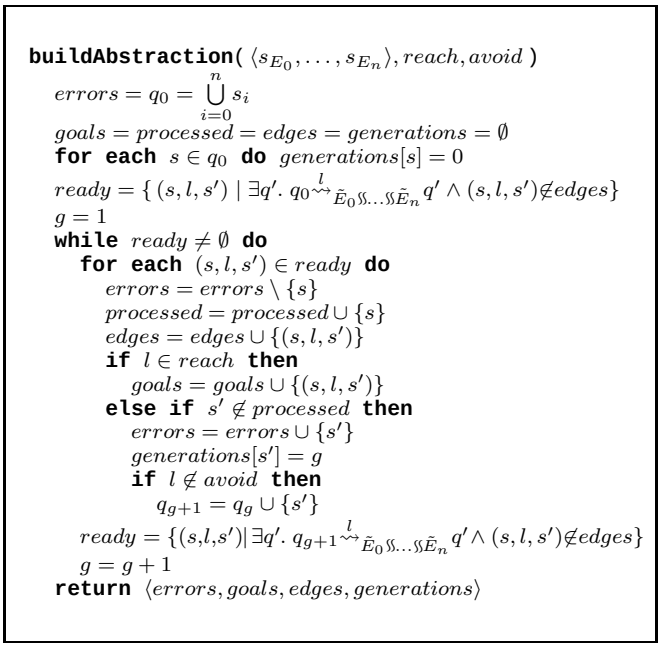
buildAbstraction($\langle s_{E_0}, \ldots, s_{E_n} \rangle, reach, avoid$)
$errors = q_0 = \bigcup_{i=0}^{n} s_i$
$goals = processed = edges = generations = \emptyset$
for each $s \in q_0$ do $generations[s] = 0$
$ready = \{ (s,l,s') \mid \exists q'.\ q_0 \overset{l}{\rightsquigarrow}_{\tilde{E}_0 \S \ldots \S \tilde{E}_n} q' \wedge (s,l,s') \notin edges \}$
$g = 1$
while $ready \neq \emptyset$ do
for each $(s,l,s') \in ready$ do
$errors = errors \setminus \{s\}$
$processed = processed \cup \{s\}$
$edges = edges \cup \{(s,l,s')\}$
if $l \in reach$ then
$goals = goals \cup \{(s,l,s')\}$
else if $s' \notin processed$ then
$errors = errors \cup \{s'\}$
$generations[s'] = g$
if $l \notin avoid$ then
$q_{g+1} = q_g \cup \{s'\}$
$ready = \{(s,l,s') \mid \exists q'.\ q_{g+1} \overset{l}{\rightsquigarrow}_{\tilde{E}_0 \S \ldots \S \tilde{E}_n} q' \wedge (s,l,s') \notin edges\}$
$g = g + 1$
return $\langle errors, goals, edges, generations \rangle$

Fig. 5. LTS abstracting composition building procedure

In Fig. 5 we present a procedure that given a state $\langle s_{E_0}, \ldots, s_{E_n} \rangle$ of the composition $E_0 \| \ldots \| E_n$, and sets of actions desired to $reach$ and to $avoid$, it returns the information required to compute a heuristic ranking:

1) $errors$, a set of states that do not conduce to a goal.
2) $goals$, a set of transitions connected by a $reach$ action.
3) $edges$ of the *abstracting path graph*.
4) $generations$ of the states in the graph.

The procedure works by iteratively building the *sequence of sets of states* produced by the abstraction (Def. 7) and, at the same time, the *abstracting path graph* (Def. 10). At each iteration a new set q_g is considered, where g indicates the current *generation*. We then compute the set of $ready$ transitions, that is, the steps available from q_g following the relaxed synchronization rule from Def. 7. Observe that once a transition is processed it is never considered again. Also note that states are considered errors until they show potential to reach a goal (i.e. deadlock states are treated as errors). Since the procedure processes each transition once its complexity is $\Theta(\left|\rightsquigarrow_{\tilde{E}_0 \S \ldots \S \tilde{E}_n}\right|)$, which considering the weakened synchronization rules, can be bigger than the actual transitions in the individual components. In the worst case, the complexity cannot surpass the connection of every state by every action, that is to say that its upper-bound complexity is $O\big((\sum_{i=0}^{n} |A_{E_i}|)(\sum_{i=0}^{n} |S_{E_i}|)^2\big)$.

Once we have built the *abstracting path graph*, obtaining a ranking for the actions enabled from the state $\langle s_A, \ldots, s_N \rangle$ is straightforward. The states reached by the $goals$ transitions are set with an estimated distance to a goal of 0, while the states in the $errors$ set are set with an ∞ estimated distance. Then, we propagate these estimates from the states to its parents following the $edges$ backwards. For each step we increase the estimated distance between parent and child according to the information stored in the $generations$ map. For a state with multiple children we keep as an estimate the minimum of its children's, since we are interested in reaching the goal as fast as possible. When this back-propagation ends, all the enabled actions have an estimated distance to the goal, thus it is just a matter of sorting them with respect to these values. However, the same action could have multiple estimated values since we allow not only intra-LTS transitions but also inter-LTS transitions; in these cases we keep the best (minimum) estimate.

Summarizing, the procedure computes – through the construction of the graph – the length of the *abstracting paths* starting at a given initial state and reaching a co-safety goal. The algorithm concentrates on the best estimates, which are used to guide the on-the-fly exploration. Thus, it effectively computes an admissible heuristic, since it can never overestimate the distance to the goal (Lemma 1).

VI. Evaluation

In this section we report on an evaluation of DCS. We compare DCS with three other approaches running on an Intel i7-3770 with 8GB of RAM:

1) Monolithic explicit state representation, previously implemented in MTSA[1] (similar to CTCT [21]).
2) Monolithic symbolic state representation using BDD, implemented in MBP [7].
3) Compositional explicit state representation implemented in Supremica (SUP) [9].

The aim of the evaluation is to assess the gains in scalability with the use of the informed search procedure. For this reason, we want case studies with the capability to scale up to higher complexities. Fortunately, both MBP and Supremica come bundled with the specification for one of the most traditional examples in controller synthesis: Line Transfer (LT); which was first introduced by Wonham [21]. The fact that the models were written by the tools authors' is important because it reduces the impact of a threat to validity that would be an ill-designed model. However, in order to be able to scale the problem up we had to extend the models, yet we maintain the same structure.

Given the complexity of: writing identical specifications in the different formal languages, finding examples that can be scaled up and presenting the results; we opt for concentrating in varying three independent parameters for the LT. This gives us a big enough number of environments with different topologies to evaluate the techniques. Still, this represents another threat to validity, since ideally the techniques should be compared with a more diverse set of inputs.

The LT consists of series of machines $M_1, M_2, \ldots, M_n$ connected by buffers $B_1, B_2, \ldots, B_n$ and ending in a special machine called Test Unit (TU). A machine M_i takes work pieces from the buffer B_{i-1} (with the exception of machine M_1 that takes the work pieces from the outside). After an undetermined amount of time, the working machine M_i outputs a processed work piece through buffer B_i. Finally, when a work piece reaches the TU it can be accepted and

[1]Available at https://bitbucket.org/dciolek/mtsa

```
Machine(Id=0) = Working[0],
  Working[w:0..W] =
    (when (w < W) get[Id]   -> Working[w+1] |
     when (w > 0) put[Id+1] -> Working[w-1] ).

TU = Idle,
  Idle    = (get[M] -> Testing ),
  Testing = (ret[1] -> reject -> Idle |
             accept -> Idle)
            +{ret[0..M]}.

Buffer(Id=0) = At[0],
  At[c:0..C] = (
    when (c > 0) get[Id] -> At[c-1] |
    when (c = 0) get[Id] -> ERROR   |
    when (c < C) put[Id] -> At[c+1] |
    when (c = C) put[Id] -> ERROR   |
    when (c < C) ret[Id] -> At[c+1] |
    when (c = C) ret[Id] -> ERROR   ).

||Plant = (forall [m:0..M-1] (
    Machine(m) || Buffer(m+1)) || TU).
```

Fig. 6. FSP model of the Transfer Line case study

taken out of the system or it can be rejected and placed back in buffer B_1 for reprocessing. The only controllable actions in this case study are the taking of work pieces. An error ensues if a machine tries to take a work piece from an empty buffer or if it tries to place a processed work piece in a full buffer. One of the goals for the controller is to *avoid* the actions that lead to errors, the other goal is to *reach* a state where a processed work piece can be accepted or rejected. We do not require the controller to achieve accepted pieces as acceptance and rejection are not decided by the controller.

The case study can be scaled in three directions:

1) $Machines$ (M): number of interconnected machines.
2) $Workload$ (W): maximum number of work pieces a machine can process simultaneously.
3) $Capacity$ (C): capacity of the buffers.

Example 5. *In Fig. 6 we present the TL model as accepted by MTSA, given in the Finite State Processes modeling language (FSP). Technically, FSP is a process calculus, in the spirit of CSP, designed to be easily machine and human readable. FSP includes standard constructs such as action prefix (`->`), external choice (`|`), alphabet extension (`+`) and parallel composition (`||`).*

We model the `Machines` *and* `TU` *separately; and afterwards the* `Buffer`*. The* `Machine` *starts idle and can* `get` *as many work pieces as its workload allows it. Uncontrollably the processing ends and the element is* `put` *in the next buffer. The* `TU` *gets an element from the last buffer and can either* `accept` *it or* `reject` *it making it* `return` *to the first buffer. Note that on attempting an invalid operation the* `Buffer` *goes into an* `ERROR` *state (i.e. a deadlock). The* `||Plant` *process represents the parallel composition between all the processes.*

In Table I we consider some "small scale" cases, in which we consider all the combinations of the parameters for the following values: M in $[4, 5, 6]$, W in $[1, 2, 3]$, C in $[1, 2, 3]$.

M	W	C	States	MTSA	DCS	MBP	SUP
4	1	1	$1.5e^4$	1.5	0.01	0.07	**0.01**
4	1	2	$3.5e^4$	5.5	0.02	0.21	**0.01**
4	1	3	$7.2e^4$	29.8	0.03	0.41	**0.03**
4	2	1	$7.7e^4$	9.1	0.01	0.11	**0.01**
4	2	2	$1.8e^5$	74.5	**0.01**	0.86	0.04
4	2	3	$3.6e^5$	OM	**0.01**	5.95	0.18
4	3	1	$2.3e^5$	OM	**0.01**	0.08	0.02
4	3	2	$5.6e^5$	OM	**0.01**	0.87	0.03
4	3	3	$2.2e^6$	OM	**0.02**	11.58	0.45
5	1	1	$1.2e^5$	33.3	0.01	0.14	**0.01**
5	1	2	$3.5e^5$	OM	**0.02**	1.43	0.09
5	1	3	$8.6e^5$	OM	**0.05**	3.53	0.72
5	2	1	$8.8e^5$	OM	**0.01**	0.25	0.09
5	2	2	$2.6e^6$	OM	**0.02**	12.27	0.62
5	2	3	$6.7e^6$	OM	**0.03**	118	62
5	3	1	$3.7e^6$	OM	**0.01**	0.26	0.22
5	3	2	$1.1e^7$	OM	**0.02**	13.28	0.43
5	3	3	$2.8e^7$	OM	**0.06**	263.72	163
6	1	1	$9.7e^5$	OM	0.22	0.47	**0.09**
6	1	2	$3.5e^6$	OM	**0.18**	14.66	0.77
6	1	3	$1.0e^7$	OM	**0.21**	37.91	TO
6	2	1	$1.1e^7$	OM	**0.02**	1.75	0.13
6	2	2	$4.0e^7$	OM	**0.07**	200	79
6	2	3	$1.2e^8$	OM	**0.25**	1397	TO
6	3	1	$5.9e^7$	OM	**0.03**	1.23	0.17
6	3	2	$2.2e^8$	OM	**0.04**	228	177
6	3	3	$6.7e^8$	OM	**0.22**	TO	TO

TABLE I

SMALL SCALE LINE TRANSFER RESULTS

We report the time in seconds required by each tool to synthesize a controller with a Time Out (TO) of 30 minutes.

As it can be seen in the table, the number of states in the environment grows very quickly and it soon takes the tools to their limits. Nonetheless, DCS is able to solve all the problems under the second mark. MTSA soon runs Out of Memory (OM) while, despite also working with a full representation of the environment, MBP solves many problems including large instances. Whereas Supremica outperforms the other methods in the simpler cases but falls short on the more complex ones.

Example 6. *Analyzing the controller generated by DCS for the TL case study, we see that the heuristic guides the search for the controller disregarding unnecessary actions such as simultaneously processing multiple work pieces. The controller found by DCS is roughly a sequence of states such that each state is connected with the following with interleaving get and put actions until an $accept$ or $reject$ becomes available. In Fig.7 we show the controller generated by DCS for 2 machines. Despite the fact that the technique pursues a reachability objective, it accidentally finds controllers with cyclic behavior, this is caused by the reopening of already explored states. However, this is not to be expected in general.*

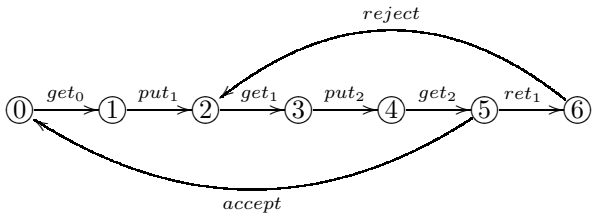

Fig. 7. Controller generated by DCS for a TL with 2 machines

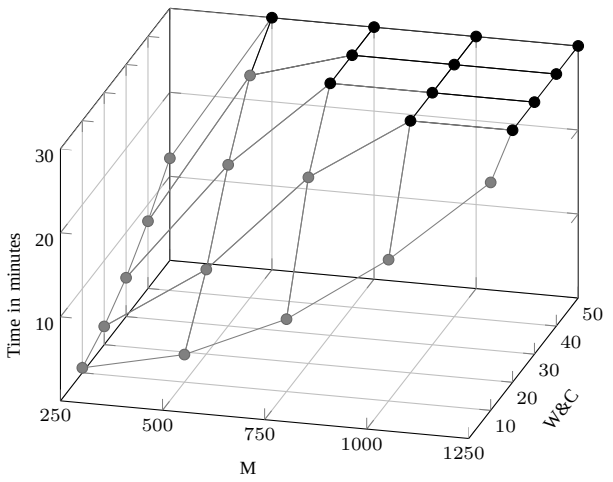

Fig. 8. Big scale LT results (TOs are depicted with black dots)

In a second experiment we assess the boundaries of the technique using "big scale" cases shown in Fig. 8. We remove the other techniques from the evaluation since they cannot tackle these problems. For ease of presentation we report on a subset of the combinations of the parameters where W and C take the same values. Note that W and C are directly related with the number of states in the Machine and Buffer LTSs respectively, while M relates to the number of components. Thus, the evaluation still allows to see how the technique performs as the complexity and the number of components increase independently.

Summarizing, we have evaluated DCS against other approaches to controller synthesis. DCS shows promise given that it is able to cope with increasing complexities beyond the capabilities of similar tools. The largest case tackled by DCS considers an environment with $8.9e^{2734}$ states (solved in $1728s$) generating a controller with 2503 states.

VII. Conclusions and future work

In this paper we present DCS, a method that looks for a controller exploring the state space on-the-fly guided by a domain-independent heuristic. The method performs better than other approaches both in time and memory for the selected case study, despite its worst case complexity. The approach forgoes maximality since reachability goals allow for a directed search in a reduced portion of the state space.

The risk of DCS is that it could fail to properly guide the exploration, triggering an unnecessary large exploration of the state space. Thus, despite our efforts to keep the approach as general as possible, further experiments are required in order to properly asses the generality of the technique. For this reason we plan to work on automatic translations to simplify the comparison with other tools.

These preliminary results show promise and hence we intend to extend the technique. Our primary concern is to deal with general liveness goals, since the domain of application of controller synthesis usually have these kind of requirements.

We highlight that a compact representation of the problem is essential for extracting useful guidance for the informed search procedure. Thus, when dealing with DES, elements such as components and synchronization provide vital information to take into account. Furthermore, even when the *LTS abstraction* achieves remarkable results, there are other heuristics that could be used (instead or in combination) that could perform better. The field of directed controller synthesis is almost uncharted, and we believe it can lead to a leap in the applicability of controller synthesis techniques.

References

[1] P. J. Ramadge and W. M. Wonham, "Supervisory Control of a Class of Discrete Event Processes," SIAM Journal on Control and Optimization, vol. 25, 1987.

[2] A. Pnueli and R. Rosner, "On the Synthesis of a Reactive Module," in Proc. 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL, 1989.

[3] J. R. Buchi and L. H. Landweber, The Collected Works of J. Richard Büchi", ch. Solving Sequential Conditions by Finite-State Strategies. Springer New York, 1990.

[4] M. O. Rabin, "Decidability of Second-order Theories and Automata on Infinite Trees," Bulletin of the American Mathematical Society, vol. 74, 1968.

[5] A. Pnueli, "The Temporal Semantics of Concurrent Programs," in Proc. of the Int. Symp. on Semantics of Concurrent Computation, 1979.

[6] Wonham, W. Murray and Ramadge, Peter J., "On the Supremal Controllable Sublanguage of a Given Language," SIAM Journal on Control and Optimization, vol. 25, no. 3, 1987.

[7] A. Gromyko, M. Pistore, and P. Traverso, "A Tool for Controller Synthesis via Symbolic Model Checking," in Proceeding of the 8th Int. Workshop on Discrete Event Systems, 2006.

[8] S. B. Akers, "Binary Decision Diagrams," IEEE Transactions in Computers, vol. 27, 1978.

[9] S. Mohajerani, R. Malik, S. Ware, and M. Fabian, "Compositional synthesis of discrete event systems using synthesis abstraction," in Control and Decision Conf., CCDC, 2011.

[10] F. Giunchiglia and P. Traverso, "Planning as Model Checking," in Recent Advances in AI Planning, vol. 1809 of Lecture Notes in Comp. Sci., 2000.

[11] B. Bonet and H. Geffner, "Planning as Heuristic Search," Artificial Intelligence, vol. 129, 2001.

[12] P. E. Hart, N. J. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," SIGART Bulletin, no. 37, 1972.

[13] C. J. Muise, S. A. McIlraith, and V. Belle, "Non-Deterministic Planning with Conditional Effects," in Proc. 24th Int. Conf. on Automated Planning and Scheduling, ICAPS, 2014.

[14] S. Edelkamp, A. Lluch-Lafuente, and S. Leue, "Directed Explicit Model Checking with HSF-SPIN," in SPIN, vol. 2057 of Lecture Notes in Comp. Sci., 2001.

[15] S. Kupferschmid, J. Hoffmann, H. Dierks, and G. Behrmann, "Adapting an AI Planning Heuristic for Directed Model Checking," in Model Checking Software, vol. 3925 of Lecture Notes in Comp. Sci., 2006.

[16] S. Tripakis and K. Altisen, "On-the-fly Controller Synthesis for Discrete and Dense Time Systems," in Proc. of World Congress on Formal Methods in the development of Computing Systems, vol. 1 of FM, 1999.

[17] V. Alimguzhin, F. Mari, I. Melatti, I. Salvo, and E. Tronci, "On-the-Fly Control Software Synthesis," in Model Checking Software, vol. 7976 of Lecture Notes in Comp. Sci., 2013.

[18] R. P. Goldman, D. J. Musliner, and M. J. S. Pelican, "Exploiting Implicit Representations in Timed Automaton Verification for Controller Synthesis," in Proc. 5th Int. Workshop of Hybrid Systems: Computation and Control, HSCC, 2002.

[19] N. D'Ippolito, D. Fischbein, M. Chechik, and S. Uchitel, "MTSA: The Modal Transition System Analyser," in Proc. 23rd IEEE/ACM Int. Conf. on Automated Software Engineering, ASE, 2008.

[20] R. Simon and R. C. T. Lee, "On the Optimal Solutions to AND/OR Series-Parallel Graphs," Journal of the ACM, vol. 18, 1971.

[21] W. Wonham, "Notes on Control of Discrete-Event Systems." Dep. of Electrical and Comp. Engineering, University of Toronto, 1999.